arXiv:1605.04729v3 [math.ST] 23 Sep 2016

Bootstrap- and Permutation-Based Inference for the Mann-Whitney Effect for Right-Censored and Tied Data

Dennis Dobler* and Markus Pauly

Ulm University, Institute of Statistics
Helmholtzstr. 20, 89081 Ulm, Germany
email: dennis.dobler@uni-ulm.de, email: markus.pauly@uni-ulm.de

September 26, 2016

Summary

The Mann-Whitney effect is an intuitive measure for discriminating two survival distributions. Here we analyze various inference techniques for this parameter in a two-sample survival setting with independent right-censoring, where the survival times are even allowed to be discretely distributed. This allows for ties in the data and requires the introduction of normalized versions of Kaplan-Meier estimators from which adequate point estimates are deduced. Asymptotically exact inference procedures based on standard normal, bootstrap- and permutation-quantiles are developed and compared in simulations. Here, the asymptotically robust and – under exchangeable data – even finitely exact permutation procedure turned out to be the best. Finally, all procedures are illustrated using a real data set.

Keywords: Counting process; Efron's bootstrap; Heteroscedasticity; Kaplan-Meier estimator; Permutation technique.

* Corresponding author

1 Introduction

When comparing the survival times from two independent groups ($j = 1, 2$) the *Mann-Whitney effect* is an intuitive measure; see e.g. Koziol and Jia (2009). In a classical survival setting with continuous life time distributions and random censoring it is given by the probability $P(T_1 > T_2)$ that a random subject from group $j = 1$ (with survival time T_1) survives longer than a randomly chosen person from group $j = 2$ (with survival time T_2). In case of uncensored data this effect reduces to the well-known Wilcoxon functional underlying the nonparametric Mann-Whitney test. Depending on the field of application it is also known as *nonparametric treatment effect* (e.g. Brunner and Munzel 2000), *stress-strength characteristic* (e.g. Kotz et al. 2003) or *probabilistic index* (e.g. Thas et al. 2012). Moreover, in case of diagnostic tests it has a direct interpretation as the area under the corresponding ROC curve, see e.g. Lange and Brunner (2012), Pauly et al. (2016) and Zapf et al. (2015) as well as Zhou et al. (2002) for more details on diagnostic accuracy measures. The Mann-Whitney effect is often estimated by the c-index for concordance (e.g. Koziol and Jia 2009). As pointed out by Acion et al. (2006), this is "*a simple, clinically relevant, and robust index*" and thus "*an ideal effect measure*", see also Kieser et al. (2013). The same still holds true in case of survival outcomes that may be subject to independent random censoring, see e.g. the glorification of the c-index in Hess (2010) or Dunkler et al. (2010). An `R`-package for a Wilcoxon-Mann-Whitney test was propagated in De Neve et al. (2014).

In the present paper we face the practically relevant situation where tied data are often inevitable. Thus, in order to take ties appropriately into account, we use a more general definition of the Mann-Whitney effect:

$$p = P(T_1 > T_2) + \frac{1}{2}P(T_1 = T_2), \tag{1.1}$$

also known as *ordinal effect size measure* in case of completely observed data (Ryu and Agresti 2008; Konietschke et al. 2012). Recently, a related effect measure, the so called *win ratio* (for prioritized outcomes), has been investigated considerably by several authors (Pocock et al. 2012, Rauch et al. 2014, Luo et al. 2015, Abdalla et al. 2016; Bebu and Lachin 2016 as well as Wang and Pocock 2016). It is the odds of the Mann-Whitney effect p in (1.1), i.e.

$$w = \frac{p}{1-p}, \tag{1.2}$$

which is also referred to as the *odds of concordance*, see Martinussen and Pipper (2013) for a treatment in the context of a semiparametric regression model. In our situation p and w describe the probability that a patient of group 1 survives longer than a patient of group 2. That is, $p > 1/2$, or equivalently $w > 1$, imply a protective survival effect for group 1. Note that until now, ties have been excluded for estimating these quantities which particularly led to the recent assessment of Wang and Pocock (2016) that "*we caution that the win ratio method should be used only when the amount of tied data is negligible*".

In this paper, we propose and rigorously study different statistical inference procedures for

both parameters p and w in a classical survival model with independent random censoring, even allowing for ties in the data. While several authors (e.g. Nandi and Aich 1994, Cramer and Kamps 1997, Kotz et al. 2003, and references therein) have considered inference for p under specific distributional assumptions, we here focus on a completely nonparametric approach, not even assuming continuity of the data. Apart from confidence intervals for p and w this also includes one- and two-sided test procedures for the null hypothesis of no group effect (tendency)

$$H_0^p : \Big\{p = \frac{1}{2}\Big\} = \Big\{w = 1\Big\}. \tag{1.3}$$

In the uncensored case this is also called the *nonparametric Behrens-Fisher problem*, see e.g. Brunner and Munzel (2000) and Neubert and Brunner (2007). To this end, the unknown parameters p and w are estimated by means of normalized versions of Kaplan-Meier estimates. These are indeed their corresponding nonparametric maximum likelihood estimates, see Efron (1967) as well as Koziol and Jia (2009) for the case of continuous observations. Based on their asymptotic properties we derive asymptotically valid tests and confidence intervals. These may be regarded as extensions of the Brunner-Munzel (2000) test to the censored data case. Since, for small sample sizes, the corresponding tests may lead to invalid α-level control (e.g. Medina et al. 2010 or Pauly et al. 2016 without censoring) we especially discuss and analyze two different resampling approaches (bootstrapping and permuting) to obtain better small sample performances.

The resulting tests are innovative in several directions compared to other existing procedures for the two-sample survival set-up:

1. We focus on the null hypothesis H_0^p in (1.3) of actual interest. Before, only the more special null hypothesis $H_0^S : \{S_1 = S_2\}$ of equal survival distributions between the two groups has been investigated, see e.g. Efron (1967), Akritas and Brunner (1997) and Akritas (2011). Corresponding one-sided testing problems (for null hypotheses formulated in terms of distribution functions) based on the related *stochastic precedence* have been treated in Arcones et al. (2002) and Davidov and Herman (2012). Instead, our procedures are not only able to assess the similarity of two survival distributions but they also quantify the degree of deviation by confidence intervals for meaningful parameters.

2. The more complex null H_0^p has so far only been studied in the uncensored case, see e.g. Janssen (1999), Brunner and Munzel (2000), De Neve et al. (2013), Chung and Romano (2016), Pauly et al. (2016) and the references cited therein. The present combination of allowing for survival analytic complications and focussing on the effect size p is achieved by utilizing empirical process theory applied to appropriate functionals.

3. We do not rely on the (elsewhere omnipresent) assumption of existing hazard rates. Instead, we here adjust for tied data by considering normalized versions of the survival function and the Kaplan-Meier estimator (leading to mid-ranks in the uncensored case). This more realistic assumption of ties in the observations accounts for a phenomenon which is oftentimes

an intrinsic problem by study design. Therefore, methodology for continuous data (even for only testing H_0^S) should not be applied. Notable exceptions for the combination of survival methods and discontinuous data are provided in Akritas and Brunner (1997) and Brendel et al. (2014) where the hypothesis H_0^S is tested.

4. Finally, small sample properties of inference procedures relying on the asymptotic theory are greatly improved by applications of resampling techniques. These utilized resampling techniques are shown to yield consistent results even in the case of ties. Thereof, the permutation procedure succeeds in being even finitely exact in the case of exchangeable survival data in both sample groups; see e.g. Lehmann and Romano (2010), Good (2010), Pesarin and Salmaso (2010, 2012), and Bonnini et al. (2014) for the classical theory on permutation tests.

In this perspective, the present paper not only states the first natural extension of point estimates for p to tied survival data, but especially introduces the first inference procedures for H_0^p (tests and confidence intervals) with rigorous consistency proofs. The latter have formerly not even been known in the continuous survival case.

The article is organized as follows. Section 2 introduces all required notation and estimators, whose combination with the variance estimator in Section 3 yields (non-resampling) inference procedures. Theoretical results concerning the resampling techniques will be presented in Sections 4.1 (the pooled bootstrap) and 4.2 (random permutation). A simulation study in Section 5 reports the improvement of the level α control by the proposed permutation and bootstrap techniques. A final application of the developed methodology to a tongue cancer data set (Klein and Moeschberger 2003) is presented in Section 6. This article's results are discussed in Section 7 and theoretically proven in Appendix A.

2 Notation, model, and estimators

For a more formal introduction of the concordance index p and the win ratio w we employ some standard notation from survival analysis. Thus, we consider two independent groups ($j = 1, 2$) of independent survival times $\tilde{T}_{j1}, \ldots, \tilde{T}_{jn_j} : (\Omega, \mathcal{A}, P) \to (0, \infty),\ j = 1, 2$, with total sample size $n = n_1 + n_2,\ n_1, n_2 \in \mathbb{N}$. Their distributions may have discrete components which reflects the situation in most clinical studies (i.e. survival times rounded to days or weeks). Since most studies pre-specify a point of time $K > 0$ after which no further observation is intended, we may also truncate the above survival times to

$$T_{ji} = \tilde{T}_{ji} \wedge K = \min(\tilde{T}_{ji}, K); \quad i = 1, \ldots, n_j,\ j = 1, 2.$$

Denote their survival functions as $S_j(t) = 1 - F_j(t) \equiv P(T_{j1} > t)$, $j = 1, 2$. Thus, both sample groups may have different, even heteroscedastic distributions. Their cumulative hazard

functions are given by $\Lambda_j(t) = -\int_0^t \frac{\mathrm{d}S_j}{S_{j-}}$ where the index minus (here in S_{j-}) always indicates the left-continuous version of a right-continuous function.

The survival times are randomly right-censored by independent, positive variables $C_{j1}, \ldots, C_{jn_j}$ with possibly discontinuous censoring survival functions $G_j, j = 1, 2$. Observation is thus restricted to

$$\mathbf{X}_j = \{(X_{j1}, \delta_{j1}), \ldots, (X_{jn_j}, \delta_{jn_j})\}, \quad j = 1, 2,$$

where $X_{ji} = \min(T_{ji}, C_{ji}), \delta_{ji} = \mathbf{1}\{X_{ji} = T_{ji}\}, 1 \leqslant i \leqslant n_j$. Note that the choice of K shall imply a positive probability of each $\{\tilde{T}_{ji} > K\}$. This constant K could, for example, be the end-of-study time, i.e. the largest censoring time. For later use we also introduce the usual counting process notation

$$\begin{aligned} N_{j;i}(u) &= \mathbf{1}\{\text{“ind. } i \text{ of group } j \text{ has an observed event during } [0, u]\text{”}\} &&= \mathbf{1}\{X_{ji} \leqslant u, \delta_{ij} = 1\}, \\ Y_{j;i}(u) &= \mathbf{1}\{\text{“ind. } i \text{ of group } j \text{ is under observation at time } u-\text{”}\} &&= \mathbf{1}\{X_{ji} \geqslant u\}. \end{aligned}$$

Summing up these quantities within each group results in $Y_j(u) = \sum_{i=1}^{n_j} Y_{j;i}(u)$, the number of group-$j$ subjects under study shortly before u, and $N_j(u) = \sum_{i=1}^{n_j} N_{j;i}(u)$, the number of observed events in group j until time u. Denote by $f^{\pm} = \frac{1}{2}(f + f_-)$ the normalized version of a right-continuous function f. With this notation the Mann-Whitney effect (1.1) and the win ratio (1.2) are given as

$$p = P(T_{11} > T_{21}) + \frac{1}{2}P(T_{11} = T_{21}) = -\int S_1^{\pm} \mathrm{d}S_2 = 1 - \int F_1^{\pm} \mathrm{d}F_2 \tag{2.1}$$

and $w = p/(1-p)$, respectively. If not specified, integration is over $[0, K]$. In this set-up we test the null hypothesis $H_0^p : \{p = 1/2\} = \{w = 1\}$ that the survival times from both groups are *tendentiously equal* against one- or two-sided alternatives. We note that the usually considered null hypothesis $H_0^S : \{S_1 = S_2\}$ of equal survival distributions is more restrictive and implies H_0^p. Similarly, a stochastic order or precedence (Davidov and Herman 2012) such as $F_1 \precnsim F_2$ implies $p \geqslant 1/2$.

In order to test H_0^S for continuous survival times, Efron (1967) has introduced a natural estimator for p, see also Koziol and Jia (2009), by replacing the unknown survival functions S_j in (2.1) with the corresponding Kaplan-Meier estimators $\widehat{S}_j, j = 1, 2$. Thus, utilizing the normalized versions of the non-parametric maximum likelihood estimators yields

$$\widehat{p} = \widehat{P}(T_{11} > T_{21}) + \frac{1}{2}\widehat{P}(T_{11} = T_{21}) = -\int \widehat{S}_1^{\pm} \mathrm{d}\widehat{S}_2 \tag{2.2}$$

and $\widehat{w} = \widehat{p}/(1-\widehat{p})$ as reasonable estimators of p and w, respectively. Similar estimators for p have been proposed by Akritas and Brunner (1997) and Brunner and Munzel (2000). The latter quantity $\widehat{w}$ has been introduced by Pocock et al. (2012) for uncensored observations (without ties) with the nice interpretation as total number of winners divided by the total number of losers in group 1 (where T_{1i} wins against $T_{2\ell}$ if $T_{1i} > T_{2\ell}$).

In order to obtain tests for H_0^p as well as one- or two-sided confidence intervals for the Mann-Whitney effect p and the win ratio w, we study the limit behaviour of $\widehat{p}$ under the general asymptotic framework

$$0 < \liminf \frac{n_1}{n} \leqslant \limsup \frac{n_1}{n} < 1 \tag{2.3}$$

as $n_1 \wedge n_2 \to \infty$. For a better illustration in intermediate steps we sometimes also assume that a unique limit exists, i.e., as $n_1 \wedge n_2 \to \infty$,

$$n_1/n \to \kappa \in (0,1). \tag{2.4}$$

Our main theorems, however, will only rely on the weak assumption of display (2.3).

We denote by "$\xrightarrow{d}$" and "$\xrightarrow{p}$" weak convergence and convergence in outer probability as $n \to \infty$, respectively, both in the sense of van der Vaart and Wellner (1996). The following central limit theorem for $\widehat{p}$ is the normalized counterpart of the asymptotics due to Efron (1967) and is proven by means of the functional δ-method in combination with the weak convergence theorem for the Kaplan-Meier estimator as stated in van der Vaart and Wellner (1996), Section 3.9.

Theorem 2.1. *(i) Suppose* (2.4) *holds, then the Mann-Whitney statistic* $V_n = V_n(p) = \sqrt{n_1 n_2/n}(\widehat{p} - p)$ *is asymptotically normal distributed, i.e.* $V_n \xrightarrow{d} Z \sim N(0, \sigma^2)$ *as* $n \to \infty$. *The limit variance is* $\sigma^2 = (1-\kappa)\sigma_{12}^2 + \kappa\sigma_{21}^2$, *where for* $j, k \in \{(1,2),(2,1)\}$*:*

$$\sigma_{jk}^2 = \int\int \Gamma_j^{\pm\pm}(u,v)\mathrm{d}S_k(u)\mathrm{d}S_k(v) \;\; \textit{and} \;\; \Gamma_j(u,v) = S_j(u)S_j(v)\int_0^{u\wedge v} \frac{\mathrm{d}\Lambda_j}{S_{j-}G_{j-}}. \tag{2.5}$$

Moreover, $\Gamma_j^{\pm\pm}$ *denotes the normalized covariance function given by*

$$\Gamma_j^{\pm\pm}(u,v) = \frac{1}{4}[\Gamma_j(u,v) + \Gamma_j(u-,v) + \Gamma_j(u,v-) + \Gamma_j(u-,v-)].$$

(ii) Under the conditions and notation of (i), the win ratio statistic $\sqrt{n_1 n_2/n}(\widehat{w} - w)$ *asymptotically follows a normal-*$N(0, \sigma^2/(1-p)^4)$*-distribution.*

Remark 2.2. *(a) Efron (1967) originally used a version of the Kaplan-Meier estimator that always considered the last observation as being uncensored. Moreover, he considered the more special*

null hypothesis $H_0^S : \{S_1 = S_2\}$ in order to simplify the variance representation under the null. (b) Without a restriction at some point of time K, a consistent estimator for the Mann-Whitney effect requires consistent estimators for the survival functions on their whole support. This involves the condition $-\int(\mathrm{d}S_j)/G_{j-} < \infty, j = 1, 2$, which is obviously only possible if the support of S_j is contained in the support of G_j; see e.g. Gill (1983), Ying (1989) and Akritas and Brunner (1997). However, this assumption is often not met in practice, e.g. if $G_j(u) = 0 < S_j(u)$ for some point of time $u > 0$.

Since the variances σ_{jk}^2 are unknown under the null H_0^p, their estimation from the data is mandatory in order to obtain asymptotically consistent tests and confidence intervals for p and w.

3 Variance estimation and studentized test statistics

Asymptotically pivotal test statistics result from studentized versions of $\widehat{p}$ and $\widehat{w}$. A natural estimator for the limit variance σ^2 of $V_n = \sqrt{n_1 n_2/n}(\widehat{p} - p)$ is $\widehat{\sigma}^2 = \frac{n_1 n_2}{n}(\widehat{\sigma}_{12}^2 + \widehat{\sigma}_{21}^2)$, where

$$n_j \widehat{\sigma}_{jk}^2 = \int\int n_j \widehat{\Gamma}_j^{\pm\pm}(u, v)\mathrm{d}\widehat{S}_k(u)\mathrm{d}\widehat{S}_k(v) \text{ and } \widehat{\Gamma}_j(u, v) = \widehat{S}_j(u)\widehat{S}_j(v)\int_0^{u\wedge v} \frac{\mathrm{d}N_j}{(1 - \frac{\Delta N_j}{Y_j})Y_j^2} \quad (3.1)$$

for $1 \leqslant j \neq k \leqslant 2$. Here, the function $\Delta f(u) = f(u) - f(u-)$ contains all jump heights of a right-continuous function f.

Lemma 3.1. *Under* (2.4) *the estimator $\widehat{\sigma}^2$ is consistent for σ^2 defined in Theorem 2.1, i.e. $\widehat{\sigma}^2 \xrightarrow{p} \sigma^2$.*

This result directly leads to the studentized statistics

$$T_n(p) = \sqrt{\frac{n_1 n_2}{n}}\frac{\widehat{p} - p}{\widehat{\sigma}} \quad \text{and} \quad W_n(w) = \sqrt{\frac{n_1 n_2}{n}}\frac{(1 - \widehat{p})^2}{\widehat{\sigma}}(\widehat{w} - w) = \sqrt{\frac{n_1 n_2}{n}}\frac{\widehat{w} - w}{\widehat{\sigma}(1 + \widehat{w})^2} \quad (3.2)$$

which are both asymptotically standard normal as $n_1 \wedge n_2 \to \infty$ by Slutzky's Theorem and the δ-method only assuming (2.3). Indeed, under (2.3), Theorem 2.1 and Lemma 3.1 might be applied along each convergent subsequence of n_1/n. Since all resulting limit distributions of $T_n(p)$ and $W_n(p)$ are pivotal, i.e. independent of κ, this weak convergence must hold for the original sequence as well. Thus, two-sided confidence intervals for p and w of asymptotic level $(1 - \alpha) \in (0, 1)$ are

$$I_n = \left[\widehat{p} \mp \frac{z_{\alpha/2}\widehat{\sigma}\sqrt{n}}{\sqrt{n_1 n_2}}\right] \text{ (for } p) \quad \text{and} \quad \left[\widehat{w} \mp \frac{z_{\alpha/2}\widehat{\sigma}}{(1 - \widehat{p})^2}\sqrt{\frac{n}{n_1 n_2}}\right] \text{ (for } w), \quad (3.3)$$

respectively, where z_α denotes the $(1-\alpha)$-quantile of the standard normal distribution. Moreover,

$$\varphi_n = \mathbf{1}\{T_n(1/2) > z_\alpha\} \quad \text{and} \quad \psi_n = \mathbf{1}\{W_n(1) > z_\alpha\} \tag{3.4}$$

are consistent asymptotic level α tests for $H_0^p : \{p = \frac{1}{2}\} = \{w = 1\}$ against the one-sided alternative hypothesis $H_1^p : \{p > \frac{1}{2}\} = \{w > 1\}$, i.e. $E(\varphi_n) \to \alpha\mathbf{1}\{p = 1/2\} + \mathbf{1}\{p > 1/2\}$ and $E(\psi_n) \to \alpha\mathbf{1}\{w = 1\} + \mathbf{1}\{w > 1\}$ as $n \to \infty$. One-sided confidence intervals and two-sided tests can be obtained by inverting the above procedures. For larger sample sizes ($n_j > 30$ depending on the magnitude of censoring), the above inference methods (3.3) and (3.4) are fairly accurate; see the simulation results in Section 5. For smaller sample sizes, however, these procedures tend to have inflated type-I error probabilities. Therefore, we propose different resampling approaches and discuss their properties in the following section. For ease of presentation, we only consider resampling tests for $H_0^p : \{p = \frac{1}{2}\}$ in order to concentrate on one parameter of interest (i.e. on p) only. Nevertheless, the results directly carry over to construct resampling-based confidence intervals for p and w, respectively.

4 Resampling the Mann-Whitney statistic

Even in the continuous case, Koziol and Jia (2009) pointed out that "with small samples sizes a bootstrap approach might be preferable" for approximating the unknown distribution of $V_n = V_n(p) = \sqrt{n_1 n_2/n}(\widehat{p} - p)$. To this end, we consider different resampling methods, starting with Efron's classical bootstrap. Here the bootstrap sample is generated by drawing with replacement from the original data pairs; see Efron (1981). Large sample properties of the bootstrapped Kaplan-Meier process and extensions thereof have been analyzed e.g. in Akritas (1986), Lo and Singh (1986) and Horvath and Yandell (1987). Calculating the bootstrap version of $\widehat{p}$ via bootstrapping for each sample group and using their quantiles leads to a slightly improved control of the type-I error probability in comparison to the asymptotic test (3.4). However, this way of bootstrapping results in a still too inaccurate behaviour in terms of too large deviations from the $\alpha = 5\%$ level (results not shown). This technique is typically improved by resampling procedures based on the pooled data $\mathbf{Z} = \{(Z_i, \eta_i) : i = 1, \ldots, n\}$ given by

$$(Z_i, \delta_i) = (X_{1i}, \delta_{1i})\mathbf{1}\{i \leqslant n_1\} + (X_{2(i-n_1)}, \delta_{2(i-n_1)})\mathbf{1}\{i > n_1\}, \quad i = 1, \ldots, n;$$

see e.g. Boos et al. (1989), Janssen and Pauls (2005) and Neubert and Brunner (2007) for empirical verifications for other functionals in this matter. Boos et al. (1989) and Konietschke and Pauly (2014) also demonstrate that random permuting of and bootstrapping from pooled samples may yield to superior results, where the first has the additional advantage of leading to finitely exact testing procedures in case of $S_1 = S_2$ and $G_1 = G_2$. We investigate both techniques in more detail below.

4.1 The pooled bootstrap

We independently draw n times with replacement from the pooled data $\mathbf{Z}$ to obtain the pooled bootstrap samples $\mathbf{Z}_1^* = (Z_{1i}^*, \eta_{1i}^*)_{i=1}^{n_1}$ and $\mathbf{Z}_2^* = (Z_{2i}^*, \eta_{2i}^*)_{i=1}^{n_2}$. Denote the corresponding Kaplan-Meier estimators based on these bootstrap samples as S_1^* and S_2^*. These may also be regarded as the n_j out of n bootstrap versions of the Kaplan-Meier estimator $\widehat{S}$ based on the pooled sample $\mathbf{Z}$. All in all, this results in the pooled bootstrap version $p^* = -\int S_1^{*\pm} \mathrm{d}S_2^*$ of $\widehat{p}$. A suitable centering term for p^* is based on the pooled Kaplan-Meier estimator and is given by $-\int \widehat{S}^{\pm} \mathrm{d}\widehat{S} = \frac{1}{2}$. Thus, we study the distribution of $V_n^* = \sqrt{\frac{n_1 n_2}{n}}(p^* - \frac{1}{2})$ for approximating the null distribution of $\sqrt{\frac{n_1 n_2}{n}}(\widehat{p} - \frac{1}{2})$.

To investigate the large sample behaviour of the bootstrap statistic p^*, first note that the pooled Kaplan-Meier estimator $\widehat{S}$ is a functional of the empirical processes based on $\mathbf{X}_1$ and $\mathbf{X}_2$. Since this functional is Hadamard-differentiable with uniformly continuous linear derivative function, Donsker theorems for the empirical processes of $\mathbf{X}_1$ and $\mathbf{X}_2$ immediately carry over to the pooled Kaplan-Meier estimator. In case of continuously distributed event times this can also be seen by utilizing the usual martingale arguments of Andersen et al. (1993), Section IV.3. In all convergence results stated below the càdlàg space $D[0, K]$ is always equipped with the $\sup$-norm; cf. van der Vaart and Wellner (1996).

Lemma 4.1. *Suppose that* (2.4) *holds. Then, as* $n \to \infty$*, we have* $\sqrt{n}(\widehat{S} - S) \xrightarrow{d} U$ *on* $D[0, K]$*, where* U *is a zero-mean Gaussian process with covariance function*

$$\Gamma : (r, s) \mapsto S(r)S(s) \int_0^{r \wedge s} \frac{\mathrm{d}\Lambda}{(1 - \Delta\Lambda)(\kappa S_{1-} G_{1-} + (1 - \kappa) S_{2-} G_{2-})}. \tag{4.1}$$

Here $\mathrm{d}\Lambda = \frac{\kappa S_{1-} G_{1-}}{\kappa S_{1-} G_{1-} + (1-\kappa) S_{2-} G_{2-}} \mathrm{d}\Lambda_1 + \frac{(1-\kappa) S_{2-} G_{2-}}{\kappa S_{1-} G_{1-} + (1-\kappa) S_{2-} G_{2-}} \mathrm{d}\Lambda_2$ *and* $S(t) = \prod_{(0,t]}(1 - \mathrm{d}\Lambda)$.

A similar behaviour of the bootstrap counterpart is established in the following theorem. Its proof relies on the δ-method for the bootstrap from which the large sample properties of V_n^* can be established as well.

Theorem 4.2. *Suppose that* (2.4) *holds. As* $n \to \infty$ *and given* $\mathbf{Z}$*, we have conditional weak convergence on* $D[0, K]$*,*

$$\sqrt{n_j}(S_j^* - \widehat{S}) \xrightarrow{d} U, \quad j = 1, 2,$$

in outer probability towards a Gaussian zero-mean process U *with covariance function* Γ *given in* (4.1)*.*

Since pooled sampling affects the covariance structure related to V_n^* (see the Appendix for details), a studentization becomes mandatory. This is also in line with the general recommendation

to bootstrap studentized statistics, see e.g. Hall and Wilson (1991), Janssen and Pauls (2005) or Delaigle et al. (2011). To this end, introduce the bootstrap variance estimator

$$\sigma^{*2} = \frac{n_2}{n}\int\int S_1^*(u)\Big[n_1\int_0^{u\wedge v}\frac{\mathrm{d}\Lambda_1^*}{(1-\Delta\Lambda_1^*)Y_1^*}\Big]S_1^*(v)\mathrm{d}S_2^*(u)S_2^*(v)$$
$$+\frac{n_1}{n}\int\int S_2^*(u)\Big[n_2\int_0^{u\wedge v}\frac{\mathrm{d}\Lambda_2^*}{(1-\Delta\Lambda_2^*)Y_2^*}\Big]S_2^*(v)\mathrm{d}S_1^*(u)S_1^*(v),$$

where N_j^* and Y_j^* are the obvious bootstrap versions of the counting processes N_j and Y_j and $\mathrm{d}\Lambda_j^* = \mathrm{d}N_j^*/Y_j^*$ define the pooled bootstrap version of the Nelson-Aalen estimator, $j = 1, 2$. We state our main result on the pooled bootstrap.

Theorem 4.3. *Suppose that* (2.3) *holds. Then the studentized bootstrap statistic* $T_n^* = V_n^*/\sigma^*$ *always approximates the null distribution of* $T_n(1/2)$ *in outer probability, i.e. we have for any choice of* p *and as* $n_1 \wedge n_2 \to \infty$*:*

$$\sup_x \left|P_p(T_n^* \leqslant x|\boldsymbol{Z}) - P_{1/2}\left(T_n(1/2) \leqslant x\right)\right| \xrightarrow{p} 0. \tag{4.2}$$

Moreover, denoting by $c_n^*(\alpha)$ *the conditional* $(1-\alpha)$*-quantile of* T_n^* *given* $\mathbf{Z}$*, it follows that* $\varphi_n^* = \mathbf{1}\{T_n(1/2) > c_n^*(\alpha)\}$ *is a consistent asymptotic level* α *test for* $H_0^p : \{p = \frac{1}{2}\}$ *against* H_1^p *that is asymptotically equivalent to* φ_n*, i.e. we have* $E(|\varphi_n^* - \varphi_n|) \to 0$*.*

4.2 Random permutation

An alternative resampling technique to Efron's bootstrap is the permutation principle. The idea is to randomly interchange the group association of all individuals while maintaining the original sample sizes. The effect measure is then calculated anew based on the permuted samples. A big advantage of permutation resampling over the pooled bootstrap is the finite exactness of inference procedures on the smaller null hypothesis

$$H_0^{S,G} : \{S_1 = S_2 \text{ and } G_1 = G_2\} \subset H_0^p; \tag{4.3}$$

see e.g. Neuhaus (1993) and Brendel et al. (2014) in case of testing H_0^S and Janssen (1997), Janssen (1999), Neubert and Brunner (2007), Chung and Romano (2013a,b), Pauly et al. (2015) as well as Pauly et al. (2016) in other situations without censoring.

Therefore, let $\pi : \Omega \to \mathcal{S}_n$ be independent of $\mathbf{Z}$ and uniformly distributed on the symmetric group $\mathcal{S}_n$, the set of all permutations of $(1, \ldots, n)$. The permuted samples are obtained as $\mathbf{Z}_1^\pi =$

$(Z_{\pi(i)}, \eta_{\pi(i)})_{i=1}^{n_1}$ and $\mathbf{Z}_2^\pi = (Z_{\pi(i)}, \eta_{\pi(i)})_{i=n_1+1}^{n}$. Plugging the Kaplan-Meier estimators S_1^π and S_2^π based on these permuted samples into the Wilcoxon functional leads to the permutation version

$$p^\pi = -\int S_1^{\pi-} \mathrm{d}S_2^\pi$$

of $\widehat{p}$. The permutation sampling is equivalent to drawing without replacement from the pooled sample $\mathbf{Z}$. The following auxiliary result for the permutation version of the Kaplan-Meier estimator is analogous to Theorem 4.2, where now $D^2[0,K]$ is equipped with the $\max$-$\sup$-norm.

Theorem 4.4. *Suppose that* (2.4) *holds. Then, as $n \to \infty$, the permutation versions of the Kaplan-Meier estimators conditionally converge on $D^2[0,K]$ in distribution*

$$(\sqrt{n_1}(S_1^\pi - \widehat{S}), \sqrt{n_2}(S_2^\pi - \widehat{S})) \xrightarrow{d} (\sqrt{1-\kappa}U, -\sqrt{\kappa}U)$$

given $\mathbf{Z}$ in outer probability. Here, U is a zero-mean Gaussian process with covariance function Γ given in (4.1).

It is shown in the Appendix, that $V_n^\pi = \sqrt{\frac{n_1 n_2}{n}}(p^\pi - \frac{1}{2})$ does in general not possess the same limit distribution as V_n. Indeed, the limit variances may be different in general and again studentization of V_n^π is necessary. This is achieved by utilizing the permutation version of $\widehat{\sigma}^2$:

$$\begin{aligned}\sigma^{\pi 2} = \frac{n_2}{n}\int\int S_1^\pi(u)\Big[n_1\int_0^{u\wedge v}\frac{\mathrm{d}N_1^\pi}{(1-\frac{\Delta N_1^\pi}{Y_1^\pi})Y_1^{\pi 2}}\Big]S_1^\pi(v)\mathrm{d}S_2^\pi(u)S_2^\pi(v)\\ +\frac{n_1}{n}\int\int S_2^\pi(u)\Big[n_2\int_0^{u\wedge v}\frac{\mathrm{d}N_2^\pi}{(1-\frac{\Delta N_2^\pi}{Y_2^\pi})Y_2^{\pi 2}}\Big]S_2^\pi(v)\mathrm{d}S_1^\pi(u)S_1^\pi(v)\end{aligned}$$

yielding the studentized permutation statistic $T_n^\pi = V_n^\pi/\sigma^\pi$. It is indeed the permutation version of $T_n(1/2)$ and also shares its asymptotic distribution as stated below.

Theorem 4.5. *Suppose that* (2.3) *holds. Then the studentized permutation statistic $T_n^\pi = V_n^\pi/\sigma^\pi$ always approximates the null distribution of $T_n(1/2)$ in outer probability, i.e. we have for any choice of p and as $n_1 \wedge n_2 \to \infty$:*

$$\sup_x \left|P_p(T_n^\pi \leqslant x|\boldsymbol{Z}) - P_{1/2}\left(T_n(1/2) \leqslant x\right)\right| \xrightarrow{p} 0. \tag{4.4}$$

Moreover, denoting by $c_n^\pi(\alpha)$ the conditional $(1-\alpha)$-quantile of T_n^π given $\mathbf{Z}$, it follows that $\varphi_n^\pi = \mathbf{1}\{T_n(1/2) > c_n^\pi(\alpha)\}$ possesses the same asymptotic properties as φ_n^* in Theorem 4.3. Furthermore,

φ_n^π is even a finitely exact level α test under $H_0^{S,G}$.

5 Finite sample properties

In this section we study the finite-sample properties of the proposed approximations. In particular, we compare the actual coverage probability of the asymptotic two-sided confidence interval I_n given in (3.3) with that of the corresponding bootstrap and permutation confidence intervals

$$I_n^* = \left[\widehat{p} \mp \frac{c_n^*(\alpha/2)\widehat{\sigma}\sqrt{n}}{\sqrt{n_1 n_2}}\right] \quad \text{and} \quad I_n^\pi = \left[\widehat{p} \mp \frac{c_n^\pi(\alpha/2)\widehat{\sigma}\sqrt{n}}{\sqrt{n_1 n_2}}\right],$$

respectively. To this end, the following distribution functions F_1 and F_2, frequently occuring in the survival context, have been chosen in our simulation study:

(1) Group 1: Exponential distribution with mean $1/2$, i.e. $F_1 = Exp(1/2)$.
Group 2: Mixture of two exponential distributions:
$F_2 = \frac{1}{3}Exp(1/1.27) + \frac{2}{3}Exp(1/2.5)$.
End-of-study time: $K \approx 1.6024$ such that $p \approx 1/2$.

(2) Group 1: Weibull distribution with scale parameter 1.65 and shape parameter 0.9.
Group 2: Standard lognormal distribution.
End-of-study time: $K \approx 1.7646$ such that $p \approx 1/2$.

(3) Groups 1 and 2: Equal Weibull distributions with scale parameters 1 and shape parameters 1.5.
End-of-study time: $K = 2$ such that $p = 1/2$.

Censoring is realized using i.i.d. exponentially distributed censoring variables C_{ji} with parameters chosen such that the (simulated) censoring probability (after truncation at K) for each of both sample groups belongs to the following ranges:

- Strong censoring: Censoring percentages between 40.97 and 43.6 per cent,
- Moderate censoring: Censoring percentages between 21.19 and 26.39 per cent,
- No censoring: Censoring percentages set to zero.

The sample sizes range over $n_1 = n_2 \in \{10, 15, 20, 25, 30\}$ as well as $n_2 = 2n_1 \in \{20, 30, 40, 50, 60\}$. Simulating 10,000 individuals each, the approximate proportions of observations greater than K are given in Table 1.

set-up / censoring	strong Group I/II	moderate Group I/II	none Group I/II
(1)	0.36 / 0.32	1.19 / 1.35	4.12 / 4.26
(2)	12.16 / 10.44	22.3 / 17.43	34.26 / 28.37
(3)	1.54	3.28	6.02

Table 1: Simulated percentages of observations greater than the respective K. In set-ups (1) and (2) the proportions for both sample groups are separated by a " / ".

The pre-specified nominal level is $1 - \alpha = 95\%$. Each simulation was carried out using $N = 10,000$ independent tests, each with $B = 1,999$ resampling steps in `R` version 3.2.3 (R Development Core Team 2016). All Kaplan-Meier estimators were calculated using the `R` package *etm* by Allignol et al. (2011). In comparison to the pooled bootstrap confidence interval I_n^*, the permutation-based confidence interval I_n^π provides even finitely exact coverage probabilities if the restricted null hypothesis $H_0^{S,G}$ given in (4.3) is true.

The simulation results for all scenarios are summarized in Tables 2 and 3. Starting with the balanced setting ($n_1 = n_2$) it can be readily seen that in case of no censoring, both, the asymptotic and permutation confidence intervals I_n and I_n^π show very accurate coverages; even for small samples.

	censoring	strong			moderate			none		
set-up	n_1	I_n	I_n^*	I_n^π	I_n	I_n^*	I_n^π	I_n	I_n^*	I_n^π
(1)	10, 10	90.63	95.70	94.95	93.49	94.63	95.08	95.02	93.80	94.63
	15, 15	92.69	95.48	94.99	94.26	94.59	95.06	94.76	93.91	94.46
	20, 20	93.08	95.83	94.65	94.00	94.20	94.63	95.56	93.84	94.98
	25, 25	93.64	95.37	94.38	94.41	94.18	94.70	94.95	93.45	94.66
	30, 30	93.96	94.91	94.31	94.73	94.12	94.27	95.34	93.80	94.37
(2)	10, 10	91.73	95.32	94.85	93.97	94.26	95.26	95.48	93.68	94.97
	15, 15	93.26	94.73	95.08	94.76	94.66	94.69	95.48	94.11	95.18
	20, 20	93.70	94.80	95.00	94.80	94.47	95.05	95.48	94.17	95.21
	25, 25	94.06	94.96	94.95	94.87	94.41	95.13	95.66	94.09	94.78
	30, 30	94.28	94.52	94.69	95.22	94.78	94.86	95.64	94.44	94.74
(3)	10, 10	89.88	95.33	94.64	92.86	94.65	94.92	94.86	94.30	94.97
	15, 15	92.05	95.60	94.82	93.85	94.58	95.09	95.57	94.06	95.28
	20, 20	92.98	95.05	95.32	94.33	94.26	95.03	95.22	94.36	95.25
	25, 25	93.32	95.12	94.78	94.07	94.32	95.11	95.34	94.38	95.04
	30, 30	93.90	95.09	94.96	94.63	94.60	95.09	94.95	94.05	94.64

Table 2: Simulated coverage probabilities (in %) of two-sided asymptotic 95% confidence intervals for $p = 0.5$ and equal sample sizes $n_1 = n_2$.

This is in line with previous findings of α-level control of rank-based tests for H_0^p (e.g. Neubert and Brunner 2007 or Pauly et al. 2016). In comparison the bootstrap based confidence intervals I_n^* are slightly liberal; especially in the exponential set-up (1) with coverages between 93.45% and 93.91%. In case of moderate and strong censoring, however, the behaviour of the asymptotic and bootstrap confidence intervals changes: The asymptotic procedure I_n based on normal quantiles gets liberal while the bootstrap procedure is more or less accurate. The liberality of I_n is particularly apparent for the smallest total sample sizes of $n = 20$ (with coverages between 89.88% and 91.73%) but also remains present with increasing sample sizes. In comparison, the permutation interval I_n^π achieved very accurate results in all set-ups. Even in the presence of strong censoring and with unequal distributions (set-ups (1) and (2)) the simulated coverage probabilities are always very close to the nominal level of 95%.

	censoring	strong			moderate			none		
set-up	n_1	I_n	I_n^*	I_n^π	I_n	I_n^*	I_n^π	I_n	I_n^*	I_n^π
(1)	10, 20	91.66	95.21	95.13	93.08	94.09	94.91	94.20	93.71	94.48
	15, 30	92.41	95.35	94.59	93.77	94.16	95.00	94.34	93.86	94.30
	20, 40	93.62	94.77	94.53	93.93	93.93	94.63	94.50	93.60	94.51
	25, 50	93.07	94.55	94.37	93.98	94.06	94.01	94.46	93.96	94.43
	30, 60	92.94	94.42	94.48	93.97	94.28	94.46	94.83	93.27	94.35
(2)	10, 20	92.02	95.21	95.11	93.43	94.49	95.09	94.93	94.33	95.03
	15, 30	92.94	95.30	95.37	94.57	94.90	95.16	95.29	94.60	95.32
	20, 40	94.01	95.45	95.29	94.45	94.75	95.19	95.46	95.67	95.98
	25, 50	93.94	95.21	94.93	95.14	94.84	95.11	95.72	95.15	95.27
	30, 60	94.24	94.97	94.93	94.73	95.30	95.00	95.74	94.81	96.05
(3)	10, 20	90.99	95.46	95.14	92.99	94.79	94.88	94.61	94.27	94.63
	15, 30	92.18	95.07	95.23	93.58	94.90	95.12	94.61	94.00	94.81
	20, 40	93.35	95.27	95.04	94.02	94.80	94.72	94.81	94.82	94.72
	25, 50	93.18	95.51	95.14	94.78	94.80	94.58	94.84	94.59	94.97
	30, 60	94.03	94.82	95.01	94.45	95.02	94.82	94.90	94.24	95.04

Table 3: Simulated coverage probabilities (in %) of two-sided asymptotic 95% confidence intervals for $p = 0.5$ and unequal sample sizes $2n_1 = n_2$.

From the results for the unbalanced case $2n_1 = n_2$, shown in Table 3, we can draw the same conclusion. In particular, the permutation approach seems to be the most promising since it does not only show the best coverages probabilities but has the additional advantage of being finitely exact in case of equal censoring and survival distributions, i.e. for $F_1 = F_2$ (implying $p = 1/2$) and $G_1 = G_2$.

All in all, the permutation procedure can be generally recommended. In the censored case the bootstrap procedure shows a similar coverage (with a minor liberality for moderate censoring) but

does not possess the nice exactness property under $H_0^{S,G}$. The asymptotic procedure can only be recommended for larger sample sizes if no or only slight censoring is apparent.

6 Application to a data example

To illustrate the practical applicability of our novel approaches we reconsider a data-set containing survival times of tongue cancer patients, cf. Klein and Moeschberger (2003). This data-set is freely available in the `R` package ***KMsurv*** via `data(tongue)`. It contains 80 patients of which $n_1 = 52$ are suffering from an aneuploid tongue cancer tumor (group 1) and $n_2 = 28$ are suffering from a diploid tumor (group 2). Observation of 21 patients in group 1 and of six patients in group 2 have been right-censored, for all others the time of death has been recorded. Thus, the corresponding censoring proportions are intermediate between the "strong" and "moderate" scenarios of Section 5. Note that the data set actually contains ties: among the uncensored survival times, there are 27 different times of death in the first group, 20 different in the second group, and 39 different in the pooled sample. There are three individuals in group 1 with censoring time exceeding the greatest recorded time of death in this group; for group 2 there is one such individual. As a reasonable value for restricting the time interval, we may thus choose $K = 200$ weeks which still precedes all just mentioned censoring times.

The Kaplan-Meier estimators correspoding to both recorded groups are plotted in Figure 1. It shows that the aneuploid Kaplan-Meier curve is always above the Kaplan-Meier curve of the diploid group. We would therefore like to examine whether this gap already yields significant results concerning the probability of concordance.

The data evaluation resulted in a point estimate $\widehat{p} \approx 0.6148$ indicating a slightly larger survival probability of the aneuploid group in comparison to the diploid. To infer this, the one- and two-sided 95% confidence intervals (based on normal, bootstrap and permutation quantiles) for the probability p that a randomly chosen individual with an aneuploid tumor survives longer than a corresponding patient with diploid tumor are given in Table 4. These have been calculated using the asymptotic normal quantile as well as $B = 9,999$ resampling iterations for each of the bootstrap and the permutation technique.

	confidence interval	
method	two-sided	one-sided
asymptotic	[0.475, 0.755]	[0.497, 1.000]
bootstrap	[0.457, 0.772]	[0.507, 1.000]
permutation	[0.464, 0.766]	[0.506, 1.000]

Table 4: One- and two-sided 95% confidence intervals based on normal, bootstrap and permutation quantiles for the Mann-Whitney effect in the tongue cancer data. In all cases the point estimate is $\widehat{p} \approx 0.6148$.

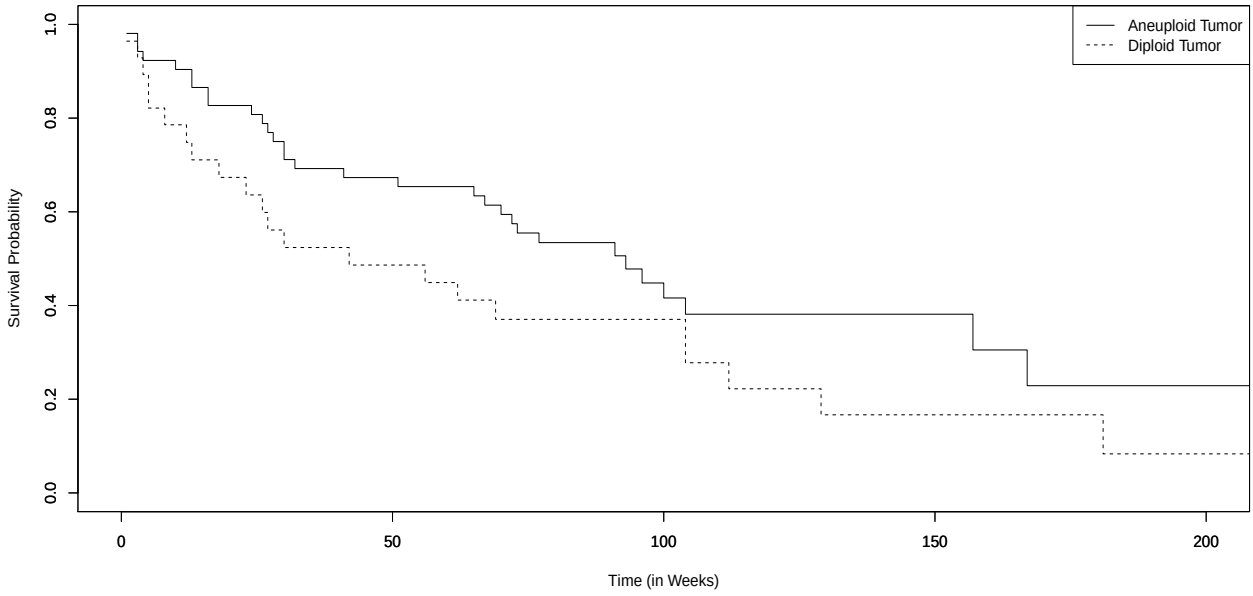

Figure 1: Kaplan-Meier estimators for patients with diploid tumor (- - -) and aneuploid tumor (——).

By inverting these confidence intervals it can be readily seen that the two-sided null hypothesis $H_0^p : \{p = 1/2\}$ cannot be rejected by any procedure since $p = 1/2$ is contained in all two-sided intervals. However, if we were only interested in detecting an effect in favor of the aneuploid group, we have to consider the corresponding one-sided tests to avoid possible directional errors. In particular, the results for testing the one-sided hypothesis $H_{0,\leqslant}^p : \{p \leqslant 1/2\}$ are borderline: It can be rejected by the resampling approaches at level 5% but lies close to the confidence limit of the asymptotic interval. Here, a slightly larger data-set might have caused a different decision. We note that multiplicity issues have not been taken into account.

7 Summary and discussion

In this article, novel inference procedures for the Mann-Whitney effect p and the win ratio w are introduced both of which are meaningful and well-established effect measures (especially in biometry and survival analysis). In comparison to the usual survival hypothesis $H_0^S : \{S_1 = S_2\}$, we were the first who particularly developed asymptotic confidence intervals for p and w as well as tests for the more interesting composite null hypothesis $H_0^p : \{p = 1/2\}$ in the two-sample survival model with right-censored data. By utilizing normalized Kaplan-Meier estimates these can even be constructed for discontinuously distributed survival times that may be subject to independent

right-censoring. Applying empirical process theory we showed that point estimates of p and w are asymptotically normal. By introducing novel variance estimates, this leads to asymptotic inference procedures based on normal quantiles. To improve their finite sample performance, bootstrap and permutation approaches have been considered and shown to maintain the same asymptotic properties. In our simulation study it could be seen that the proposed permutation procedure considerably improves the finite sample performance of our procedure. Moreover, it is even finitely exact if data is exchangeable (i.e., whenever both survival and both censoring distributions are equal) and can thus be recommended as the method of choice. In the special continuous situation with complete observations a similar result has been recently proven in Chung and Romano (2016).

Note, that the proposed method can also be applied in the 'winner-loser' set-ups considered in Pocock et al. (2012) or Wang and Pocock (2016), where now even the neglected ties can be taken into account. We plan to do this in the near future. Moreover, extensions of the proposed techniques to other models such as multiple samples and multivariate or specific paired designs (e.g. measurements before and after treatment) will also be considered in a forthcoming paper.

Appendix

A Proofs

Proof of Theorem 2.1: Integration by parts shows that

$$\begin{aligned} p &= -\int S_1^{\pm} \mathrm{d}S_2 = -\frac{1}{2}\int S_1 \mathrm{d}S_2 + \frac{1}{2}\int S_2 \mathrm{d}S_1 - \frac{1}{2}S_1(K)S_2(K) + \frac{1}{2}S_1(0)S_2(0) \\ &= -\frac{1}{2}\int S_1 \mathrm{d}S_2 + \frac{1}{2}\int S_2 \mathrm{d}S_1 + \frac{1}{2} \\ &= -\frac{1}{2}\int_{[0,K)} S_1 \mathrm{d}S_2 + \frac{1}{2}\int_{[0,K)} S_2 \mathrm{d}S_1 + \frac{1}{2} =: \phi(S_1, S_2). \end{aligned}$$

Here it is important to exclude the right boundary K of the support of S_1 and S_2 since otherwise theorems on the weak convergence of the Kaplan-Meier estimator on the whole support of the survival function would be required; see Gill (1983) and Ying (1989) for such statements as well as Dobler (2016) for a bootstrap version. Each integral is a Hadamard-differentiable functional of (S_1, S_2) tangentially to

$$\begin{aligned} \mathcal{H}^2 := \{(h_1, h_2) \in D^2[0,K] : h_1^{\pm} \in \mathcal{L}_1|_{[0,K]}(S_2), h_2^{\pm} \in \mathcal{L}_1|_{[0,K]}(S_1), \\ h_1(0) = h_2(0), h_1(K) = h_2(K) = 0\}; \end{aligned}$$

see e.g. van der Vaart and Wellner (1996), Lemma 3.9.17. Note that S_1, S_2 are monotone and thus of bounded variation. The continuous, linear Hadamard-derivative is given by

$$\mathrm{d}\phi_{(S_1,S_2)} \cdot (h_1, h_2) = \frac{1}{2}\Big[- \int_{[0,K)} S_1 \mathrm{d}h_2 - \int_{[0,K)} h_1 \mathrm{d}S_2 + \int_{[0,K)} S_2 \mathrm{d}h_1 + \int_{[0,K)} h_2 \mathrm{d}S_1 \Big],$$

integrals with respect to h_1, h_2 defined via integration by parts. Because the integrands and integrators have no mass in K, this derivative can be further simplified to

$$\mathrm{d}\phi_{(S_1,S_2)} \cdot (h_1, h_2) = - \int h_1^{\pm} \mathrm{d}S_2 + \int h_2^{\pm} \mathrm{d}S_1, \tag{A.1}$$

again via integration by parts. By the functional δ-method (Section 3.9 in van der Vaart and Wellner 1996), we thus conclude that

$$\begin{aligned}
V_n(p) &= \sqrt{\frac{n_1 n_2}{n}}[\phi(\widehat{S}_1, \widehat{S}_2) - \phi(S_1, S_2)] \\
&= \mathrm{d}\phi_{(S_1,S_2)} \cdot \sqrt{\frac{n_1 n_2}{n}}(\widehat{S}_1 - S_1, \widehat{S}_2 - S_2) + o_p(1) \\
&= \sqrt{\frac{n_1}{n}} \int \sqrt{n_2}(\widehat{S}_2 - S_2)^{\pm} \mathrm{d}S_1 - \sqrt{\frac{n_2}{n}} \int \sqrt{n_1}(\widehat{S}_1 - S_1)^{\pm} \mathrm{d}S_2 + o_p(1) \\
&\xrightarrow{d} \sqrt{\kappa} \int U_2^{\pm} \mathrm{d}S_1 + \sqrt{1-\kappa} \int U_1^{\pm} \mathrm{d}S_2 =: Z
\end{aligned}$$

as $n \to \infty$, where U_1 and U_2 are independent, $D[0, K]$-valued, zero-mean Gaussian processes with covariance functions

$$\Gamma_j : (r, s) \mapsto 1\{r, s < K\} S_j(r) S_j(s) \int_0^{r \wedge s} \frac{\mathrm{d}\Lambda_j}{(1 - \Delta\Lambda_j) S_{j-} G_{j-}}, \quad j = 1, 2,\ 0 \leqslant r, s \leqslant K$$

and with $U_1(K) = U_2(K) \equiv 0$; see Example 3.9.31 in van der Vaart and Wellner (1996). Thus, $E(Z) = 0$ and

$$\sigma^2 = var(Z) = \kappa \int \int \Gamma_2^{\pm\pm} \mathrm{d}S_1 \mathrm{d}S_1 + (1 - \kappa) \int \int \Gamma_1^{\pm\pm} \mathrm{d}S_2 \mathrm{d}S_2$$

where again $\Gamma_j^{\pm\pm}(r, s) = \frac{1}{4}[\Gamma_j(r, s) + \Gamma_j(r, s-) + \Gamma_j(r-, s) + \Gamma_j(r-, s-)],\ j = 1, 2.$ □

Proof of Lemma 3.1: The consistency $\widehat{\sigma}^2 \xrightarrow{p} \sigma^2$ follows from the consistency of the Kaplan-Meier estimator, the continuity of $D^2[0, K] \ni (f, g) \mapsto \int f \mathrm{d}g$ in all functions g of bounded variation (see the proof of Lemma 3 in Gill 1989), as well as from the continuity of $D[0, K] \ni f \mapsto \frac{1}{f}$ in all

functions that are bounded away from zero. □

Proof of Lemma 4.1: This statement is proven basically in the same way as the large-sample properties of the sample-specific Kaplan-Meier estimators, i.e., applying the functional δ-method (Theorem 3.9.5 in van der Vaart and Wellner 1996) to the (pooled) empirical process of $\mathbf{Z}$, indexed by $\mathcal{F} = \{1\{\,\cdot_1 \leqslant z,\, \cdot_2 = 1\},\, 1\{\,\cdot_1 \geqslant z\} : \, z \in [0, K)\}$. The only difference is the limit of the pooled Nelson-Aalen estimator: Writing $N = N_1 + N_2$ and $Y = Y_1 + Y_2$ and letting $t \in [0, K]$, we have

$$\begin{aligned}\widehat{\Lambda}(t) &= \int_0^t \frac{\mathrm{d}N}{Y} = \frac{n_1}{n}\int_0^t \frac{Y_1}{n_1}\frac{n}{Y}\frac{\mathrm{d}N_1}{Y_1} + \frac{n_2}{n}\int_0^t \frac{Y_2}{n_2}\frac{n}{Y}\frac{\mathrm{d}N_2}{Y_2} \\ &\xrightarrow{p} \kappa\int_0^t \frac{S_{1-}G_{1-}}{\kappa S_{1-}G_{1-} + (1-\kappa)S_{2-}G_{2-}}\mathrm{d}\Lambda_1 + (1-\kappa)\int_0^t \frac{S_{2-}G_{2-}}{\kappa S_{1-}G_{1-} + (1-\kappa)S_{2-}G_{2-}}\mathrm{d}\Lambda_2.\end{aligned}$$

Thus, substituting this quantity for the cumulative hazard function in the asymptotic covariance function of the Kaplan-Meier estimator, the proof is complete. □

Proof of Theorem 4.2: This is simply an application of the δ-method for the bootstrap (Theorem 3.9.11 in van der Vaart and Wellner 1996) to the result of Lemma 4.1 and the two-sample bootstrap Donsker Theorem 3.7.6 in van der Vaart and Wellner (1996). Note, that we employ a Donsker theorem for

$$\sqrt{\frac{n_1 n_2}{n}}\Big(\mathbb{P}_n^{\mathbf{Z}} - \frac{n_1}{n}P^{(X_1,\delta_1)} - \frac{n_2}{n}P^{(X_n,\delta_n)}\Big),$$

$\mathbb{P}_n^{\mathbf{Z}}$ being the empirical process of the pooled sample $\mathbf{Z}$. This yields the same limit distribution of $\sqrt{n_j}(S_j^* - \widehat{S})$ as in Lemma 4.1 by the uniform Hadamard-differentiability of all involved functionals, concluding the proof. □

Proof of Theorem 4.3: This is just another application of the functional δ-method for the bootstrap, applied to the Wilcoxon statistic ϕ and the intermediate Theorem 4.2. To see the consistency of the bootstrap variance estimator, combine the Glivenko-Cantelli theorems for the bootstrap empirical processes of $\mathbf{Z}_1^*$ and $\mathbf{Z}_2^*$ with the continuous mapping theorem. Finally, apply Slutsky's lemma; e.g. Example 1.4.7 in van der Vaart and Wellner (1996). □

Proof of Theorem 4.4: This follows in the same way as Theorem 4.2, but with the asymptotic covariance function determined by the Donsker theorem for the permutation empirical process (cf. Theorem 3.7.1 in van der Vaart and Wellner 1996) and the observation that

$$\mathbb{P}_{n_1}^{\mathbf{X}_1} - \mathbb{P}_n^{\mathbf{Z}} \;=\; \frac{n_2}{n}(\mathbb{P}_{n_1}^{\mathbf{X}_1} - \mathbb{P}_{n_2}^{\mathbf{X}_2}) \;=\; -\frac{n_2}{n_1}(\mathbb{P}_{n_2}^{\mathbf{X}_2} - \mathbb{P}_n^{\mathbf{Z}}).$$

Here $\mathbb{P}_{n_1}^{\mathbf{X}_1}$ and $\mathbb{P}_{n_2}^{\mathbf{X}_2}$ are the empirical processes based on $\mathbf{X}_1$ and $\mathbf{X}_2$, respectively; cf. p. 361 in van der Vaart and Wellner (1996). □

Proof of Theorem 4.5: This theorem is proven in the same way as Theorem 4.3, using the intermediate result of Theorem 4.4. However, the asymptotic variance of V_n^π requires some attention: By

the functional δ-method V_n^π is asymptotically equivalent to the image of

$$\sqrt{\frac{n_1 n_2}{n}}(\mathbb{P}_{n_1}^{\mathbf{Z}_1^\pi} - \mathbb{P}_n^{\mathbf{Z}}, \mathbb{P}_{n_2}^{\mathbf{Z}_2^\pi} - \mathbb{P}_n^{\mathbf{Z}}) \xrightarrow{d} ((1-\kappa)U, -\kappa U)$$

under a continuous linear map, where $\mathbb{P}_{n_j}^{\mathbf{Z}_j^\pi}$ is the jth permutation empirical process, $j = 1, 2$, and the Gaussian process U is given by Theorem 4.4. But this map (i.e., the Hadamard derivative) also subtracts both components of the previous display, cf. (A.1). Hence, T_n^π has the same large sample behaviour as its bootstrap counterpart T_n^* of Theorem 4.3. The finite exactness under exchangeable data is a well-known property of permutation tests, see e.g. Janssen and Völker (2007). □

References

S. Abdalla, M. E. Montez-Rath, P. S. Parfrey, and G. M. Chertow. The win ratio approach to analyzing composite outcomes: An application to the EVOLVE trial. *Contemporary Clinical Trials*, 48:119–124, 2016.

L. Acion, J. J. Peterson, S. Temple, and S. Arndt. Probabilistic index: An intuitive non-parametric approach to measuring the size of treatment effects. *Statistics in Medicine*, 25(4):591–602, 2006.

M. G. Akritas. Bootstrapping the Kaplan-Meier Estimator. *Journal of the American Statistical Association*, 81(396):1032–1038, 1986.

M. G. Akritas. Nonparametric Models for ANOVA and ANCOVA Designs. In *International Encyclopedia of Statistical Science*, pages 964–968. Springer, 2011.

M. G. Akritas and E. Brunner. Nonparametric Methods for Factorial Designs with Censored Data. *Journal of the American Statistical Association*, 92(438):568–576, 1997.

A. Allignol, M. Schumacher, and J. Beyersmann. Empirical Transition Matrix of Multi-State Models: The etm Package. *Journal of Statistical Software*, 38(4):1–15, 2011. URL `http://www.jstatsoft.org/v38/i04/`.

P. K. Andersen, Ø. Borgan, R. D. Gill, and N. Keiding. *Statistical Models Based on Counting Processes*. Springer, New York, 1993.

M. A. Arcones, P. H. Kvam, and F. J. Samaniego. Nonparametric Estimation of a Distribution Subject to a Stochastic Precedence Constraint. *Journal of the American Statistical Association*, 97(457):170–182, 2002.

I. Bebu and J. M. Lachin. Large sample inference for a win ratio analysis of a composite outcome based on prioritized components. *Biostatistics*, 17(1):178–187, 2016.

S. Bonnini, L. Corain, M. Marozzi, and L. Salmaso. *Nonparametric Hypothesis Testing: Rank and Permutation Methods with Applications in R*. John Wiley & Sons, 2014.

D. Boos, P. Janssen, and N. Veraverbeke. Resampling from centered data in the two-sample problem. *Journal of Statistical Planning and Inference*, 21(3):327–345, 1989.

M. Brendel, A. Janssen, C.-D. Mayer, and M. Pauly. Weighted logrank permutation tests for randomly right censored life science data. *Scandinavian Journal of Statistics*, 41(3):742–761, 2014.

E. Brunner and U. Munzel. The Nonparametric Behrens-Fisher Problem: Asymptotic Theory and a Small-Sample Approximation. *Biometrical Journal*, 42(1):17–25, 2000.

E. Chung and J. P. Romano. Exact and asymptotically robust permutation tests. *The Annals of Statistics*, 41(2):484–507, 2013a.

E. Chung and J. P. Romano. Multivariate and multiple permutation tests. Technical report, Dept. Statistics, Stanford University, 2013b.

E. Chung and J. P. Romano. Asymptotically valid and exact permutation tests based on two-sample U-statistics. *Journal of Statistical Planning and Inference*, 168:97–105, 2016.

E. Cramer and U. Kamps. The UMVUE of $P(X < Y)$ based on type-II censored samples from Weinman multivariate exponential distributions. *Metrika*, 46(1):93–121, 1997.

O. Davidov and A. Herman. Ordinal dominance curve based inference for stochastically ordered distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(5): 825–847, 2012.

O. Davidov and S. Peddada. The linear stochastic order and directed inference for multivariate ordered distributions. *Annals of Statistics*, 41(1):1–40, 2013.

J. De Neve, O. Thas, J.-P. Ottoy, and L. Clement. An extension of the Wilcoxon-Mann-Whitney test for analyzing RT-qPCR data. *Statistical Applications in Genetics and Molecular Biology*, 12(3):333–346, 2013.

J. De Neve, J. Meys, J.-P. Ottoy, L. Clement, and O. Thas. unifiedWMWqPCR: the unified Wilcoxon–Mann–Whitney test for analyzing RT-qPCR data in R. *Bioinformatics*, 30(17):2494–2495, 2014.

A. Delaigle, P. Hall, and J. Jin. Robustness and accuracy of methods for high dimensional data analysis based on student's t-statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(3):283–301, 2011.

D. Dobler. Bootstrapping the Kaplan-Meier Estimator on the Whole Line. *Preprint arXiv:1507.02838*, 2016.

D. Dunkler, M. Schemper, and G. Heinze. Gene selection in microarray survival studies under possibly non-proportional hazards. *Bioinformatics*, 26(6):784–790, 2010.

B. Efron. The two sample problem with censored data. *In Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, 4:831–853, 1967.

B. Efron. Censored Data and the Bootstrap. *Journal of the American Statistical Association*, 76 (374):312–319, 1981.

R. D. Gill. Large Sample Behaviour of the Product-Limit Estimator on the Whole Line. *The Annals of Statistics*, 11(1):49–58, 1983.

R. D. Gill. Non- and Semi-Parametric Maximum Likelihood Estimators and the von Mises Method (Part 1)[with Discussion and Reply]. *Scandinavian Journal of Statistics*, 16(2):97–128, 1989.

P. I. Good. *Permutation Tests: A Practical Guide to Resampling Methods for Testing Hypotheses*. John Wiley & Sons, New York, second edition, 2010.

P. Hall and S.R. Wilson. Two guidelines for bootstrap hypothesis testing. *Biometrics*, 47(2): 757–762, 1991.

K. R. Hess. Comparing survival curves using an easy to interpret statistic. *Clinical Cancer Research*, 16(20):4912–4913, 2010.

L. Horvath and B. Yandell. Convergence Rates for the Bootstrapped Product-Limit Process. *The Annals of Statistics*, 15(3):1155–1173, 1987.

A. Janssen. Studentized permutation tests for non-i.i.d. hypotheses and the generalized Behrens-Fisher problem. *Statistics & Probability Letters*, 36(1):9–21, 1997.

A. Janssen. Testing nonparametric statistical functionals with applications to rank tests. *Journal of Statistical Planning and Inference*, 81(1):71–93, 1999.

A. Janssen and T. Pauls. A Monte Carlo comparison of studentized bootstrap and permutation tests for heteroscedastic two-sample problems. *Computational Statistics*, 20(3):369–383, 2005.

A. Janssen and D. Völker. Most powerful conditional tests. *Statistics & Decisions*, 25(1/2007): 41–62, 2007.

M. Kieser, T. Friede, and M. Gondan. Assessment of statistical significance and clinical relevance. *Statistics in Medicine*, 32(10):1707–1719, 2013.

J. P. Klein and M. L. Moeschberger. *Survival Analysis: Techniques for Censored and Truncated Data*. Springer Science & Business Media, 2003.

F. Konietschke and M. Pauly. Bootstrapping and permuting paired t-test type statistics. *Statistics and Computing*, 24(3):283–296, 2014.

F. Konietschke, L. A. Hothorn, and E. Brunner. Rank-based multiple test procedures and simultaneous confidence intervals. *Electronic Journal of Statistics*, 6:738–759, 2012.

S. Kotz, Y. Lumelskii, and M. Pensky. *The stress–strength model and its generalizations: theory and applications*. World Scientific, Singapore, 2003.

J. A. Koziol and Z. Jia. The concordance index C and the Mann–Whitney parameter $Pr(X > Y)$ with randomly censored data. *Biometrical Journal*, 51(3):467–474, 2009.

K. Lange and E. Brunner. Sensitivity, specificity and ROC-curves in multiple reader diagnostic trialsa unified, nonparametric approach. *Statistical Methodology*, 9(4):490–500, 2012.

E. L. Lehmann and J. P. Romano. *Testing Statistical Hypotheses*. Springer Science & Business Media, New York, third edition, 2010.

S.-H. Lo and K. Singh. The Product-Limit Estimator and the Bootstrap: Some Asymptotic Representations. *Probability Theory and Related Fields*, 71(3):455–465, 1986.

X. Luo, H. Tian, S. Mohanty, and W. Y. Tsai. An Alternative Approach to Confidence Interval Estimation for the Win Ratio Statistic. *Biometrics*, 71(1):139–145, 2015.

T. Martinussen and C. B. Pipper. Estimation of odds of concordance based on the Aalen additive model. *Lifetime Data Analysis*, 19(1):100–116, 2013.

J. Medina, D. Y. Kimberg, A. Chatterjee, and H. B. Coslett. Inappropriate usage of the Brunner–Munzel test in recent voxel-based lesion-symptom mapping studies. *Neuropsychologia*, 48(1): 341–343, 2010.

S. B. Nandi and A. B. Aich. A note on confidence bounds for $P(X > Y)$ in bivariate normal samples. *Sankhyā: The Indian Journal of Statistics, Series B*, 56(2):129–136, 1994.

K. Neubert and E. Brunner. A studentized permutation test for the non-parametric Behrens–Fisher problem. *Computational Statistics & Data Analysis*, 51(10):5192–5204, 2007.

G. Neuhaus. Conditional rank tests for the two-sample problem under random censorship. *The Annals of Statistics*, 21(4):1760–1779, 1993.

M. Pauly, E. Brunner, and F. Konietschke. Asymptotic permutation tests in general factorial designs. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 77(2):461–473, 2015.

M. Pauly, T. Asendorf, and F. Konietschke. Permutation tests and confidence intervals for the area under the ROC-Curve. *Biometrical Journal, to appear*, 2016.

F. Pesarin and L. Salmaso. *Permutation tests for complex data: theory, applications and software*. John Wiley & Sons, 2010.

F. Pesarin and L. Salmaso. A review and some new results on permutation testing for multivariate problems. *Statistics and Computing*, 22(2):639–646, 2012.

S. J. Pocock, C. A. Ariti, T. J. Collier, and D. Wang. The win ratio: a new approach to the analysis of composite endpoints in clinical trials based on clinical priorities. *European heart journal*, 33 (2):176–182, 2012.

R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016. URL `http://www.R-project.org`.

G. Rauch, A. Jahn-Eimermacher, W. Brannath, and M. Kieser. Opportunities and challenges of combined effect measures based on prioritized outcomes. *Statistics in Medicine*, 33(7):1104–1120, 2014.

E. Ryu and A. Agresti. Modeling and inference for an ordinal effect size measure. *Statistics in Medicine*, 27(10):1703–1717, 2008.

O. Thas, J. De Neve, L. Clement, and J.-P. Ottoy. Probabilistic index models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(4):623–671, 2012.

A. W. van der Vaart and J. Wellner. *Weak Convergence and Empirical Processes*. Springer, New York et al., 1996.

D. Wang and S. Pocock. A win ratio approach to comparing continuous non-normal outcomes in clinical trials. *Pharmaceutical Statistics*, 15(3):238–245, 2016.

Z. Ying. A note on the asymptotic properties of the product-limit estimator on the whole line. *Statistics & Probability Letters*, 7(4):311–314, 1989.

A. Zapf, E. Brunner, and F. Konietschke. A wild bootstrap approach for the selection of biomarkers in early diagnostic trials. *BMC Medical Research Methodology*, 15(1):43, 2015.

X.-H. Zhou, D. K. McClish, and N. A. Obuchowski. *Statistical Methods in Diagnostic Medicine*. John Wiley & Sons, New York, 2002.